Rapid Fitting of Band-Excitation Piezoresponse Force Microscopy Using Physics Constrained Unsupervised Neural Networks

Alibek T. Kaliyev
Lehigh University
Bethlehem, Pennsylvania
alk224@lehigh.edu

Ryan Forelli
Lehigh University
Bethlehem, Pennsylvania
rff224@lehigh.edu

Shuyu Qin
Lehigh University
Bethlehem, Pennsylvania
shq219@lehigh.edu

Yichen Guo
Lehigh University
Bethlehem, Pennsylvania
yig319@lehigh.edu

Pedro Sales
Massachusetts Institute of Technology (MIT)
Cambridge, Massachusetts
psales@mit.edu

Seda Ogrenci-Memik
Northwestern University
Evanston, Illinois
seda@northwestern.edu

Michael W. Mahoney
University of California
Berkeley, Berkeley, California
mmahoney@stat.berkeley.edu

Amir Gholami
University of California
Berkeley, Berkeley, California
amirgh@berkeley.edu

Rama K. Vasudevan
Oak Ridge National Laboratory
Oak Ridge, Tennessee
vasudevanrk@ornl.gov

Stephen Jesse
Oak Ridge National Laboratory
Oak Ridge, Tennessee
sjesse@ornl.gov

Nhan Tran
Fermi National Accelerator Laboratory
Batavia, Illinois
ntran@fnal.gov

Philip Harris
Massachusetts Institute of Technology (MIT)
Cambridge, Massachusetts
pcharris@mit.edu

Martin Takáč
Mohamed bin Zayed University
of Artificial Intelligence
Abu Dhabi, United Arab Emirates
martin.takac@mbzuai.ac.ae

Joshua C. Agar
Drexel University
Philadelphia, Pennsylvania
jca92@drexel.edu

Abstract

Scanning probe spectroscopy generates high-dimensional data that is difficult to analyze in real-time, hindering researcher creativity. Machine learning techniques like PCA accelerate analysis but are inefficient, noise-sensitive, and lack interpretability. We developed an unsupervised deep neural network constrained by

37th Conference on Neural Information Processing Systems (NeurIPS 2023).

a known empirical equation to enable real-time, robust fitting. Demonstrated on band-excitation piezoresponse force microscopy, our model fits cantilever response to a simple harmonic oscillator more than 4 orders of magnitude faster than least squares while enhancing robustness. It performs well on noisy data where conventional methods fail. Quantization-aware training enables sub-millisecond streaming inference on a field-programmable gate array (FPGA), orders of magnitude faster than data acquisition. This methodology broadly applies to spectroscopic fitting and provides a pathway for real-time control and interpretation.

1 Introduction

Scientific discovery relies on experiments to observe, validate, and model phenomena, including nanoscale dynamics in catalysis[1], energy conversion[2], and next-generation electronics[3]. Advances in instrumentation [4][5] have vastly increased data volume, but most of it remains under-analyzed, leading to failed experiments. Analyzing data is time-consuming, limiting experiment quantity and complexity.

Recent trends in multimodal scanning and electron microscopy capture dynamic material responses[6]. Electron microscopy can reach speeds of 400 Gbps[7] while scanning probe spectroscopies provide insights into various material dynamics [8][9][10]. However, analyzing the large, high-dimensional data often occurs long after experiments, hindering real-time feedback for sensitive samples and creative inquiry (*see Supplemental Materials S1 for more information*).

2 Related work

2.1 Limitations of prior work

Machine learning (ML) holds the potential for expediting microscopy analysis but faces challenges. Traditional clustering struggles with continuous features and material physics boundaries [11][12]. Statistical dimensionality reduction[13], like PCA, is fast but less effective in multimodal materials microscopy[14][15], lacking scalability and interpretability. Other ML methods, such as dictionary learning[16] and non-negative matrix factorization[17], have their limitations.

Recent advances in deep neural networks (DNNs) with autoencoder structures enable microscopy dimensionality reduction[9][18][19][20][21]. These networks use convolutions or recurrence to capture spatial information and consider data's ordinal nature. Regularized bottleneck layers (e.g., L1 sparsity[9], variational constraints[18], or normalizing flows[22]) facilitate descriptive latent spaces, although sometimes non-physics-compliant.

Physics-informed neural networks are emerging, incorporating physics constraints (e.g., partial differential equations) into the loss function[23]. These constraints can be soft[24] or hard[25], or include empirical expressions[26], enabling interpretable data parameterizations[27]. However, applying them to real noisy materials data requires further exploration.

Integrating neural networks as fast physics-constrained approximations holds promise for real-time physics-compliant analysis in materials synthesis, characterization, and manufacturing.

2.2 Experimental procedure

Thin films of $\mathrm{PbZr0.2Ti0.8O3/Ba0.5Sr_{0.5}RuO_3/NdScO_3}$ were grown via pulsed laser deposition[9][15][18]. These films exhibit a unique hierarchical domain structure[28] due to strain from $\mathrm{NdScO_3}$, leading to variations in piezoresponse and switching mechanisms.

Band-excitation piezoresponse force microscopy (BE-PFM) was conducted on these films. BE-PFM Polarization Spectroscopy was performed in a 60×60 grid using a conductive cantilever-mounted tip, applying a small-signal BE waveform (1 V) near the cantilever resonance. Real and imaginary components of the cantilever response were obtained via fast-Fourier transform. Bipolar-triangular-switching waveforms were applied during measurements. Piezoresponse parameters (amplitude (A), phase (ϕ), resonance frequency (ω), and dissipation (q) were determined by fitting the cantilever

response to a simple harmonic oscillator model. Piezoelectric hysteresis loops were calculated for each pixel (more information in *Supplemental Materials S1*).

Conventional data analysis methods have challenges, like Least Square Fitting (LSQF[29]). They require initial parameter guesses, are sensitive to noise, and can get stuck in local minima. Signal processing and machine learning can help but need careful tuning. A neural network for initial guesses requires labeled training data and doesn't significantly speed up analysis [30].

LSQF typically takes around 15 minutes for a dataset of 1.38 million fits parallelized on 8 CPU threads. Increasing computing resources poses practical challenges. Thus, analysis often occurs well after experiments, limiting creative inquiry and the full potential of automated techniques.

3 Results and discussion

Here, we develop an unsupervised physics-constrained neural network to accelerate and robustly conduct SHO and piezoelectric hysteresis loops in experimental data with low SnR. To do this, we built an ultracompact unsupervised DNN with 7254 parameters to conduct fast approximates of fits to empirical physics-based models. With strong physics regularizers, we went beyond first-order optimization methods to improve optimization performance and speed by using approximate second-order optimization methods.

Overall, we demonstrate several important breakthroughs:

1. **Speed** – we can train our model to conduct 1.38 million SHO fits in less than 5 minutes and can conduct inference in <3 seconds with a batch size of 1024 on free computing resources (PCIe P100 on Google Colab).
2. **Robustness** – We demonstrate that SHO and hysteresis loop fit results have narrower and more physically reasonable distributions than LSQF results.
3. **Signal-to-noise** – Our model performs well and provides physically interpretable results on artificially noisy data where well-designed conventional LSQF pipelines fail.
4. **Real-time** – We conduct quantization-aware training to deploy this model on an FPGA. Simulations predict streaming inference at 37 μs, orders of magnitude faster than the data acquisition and sufficiently fast for real-time control of automated experiments.

To conduct SHO and piezoelectric hysteresis loop fits, we designed a physics-constrained DNN (Figure 1). The general philosophy is to design an autoencoder that takes the signal, the real and imaginary components, as input. It predicts an embedding (x), where x are the parameters of the empirical fitting function. Thus, the decoder is merely the empirical function shown in the inset of Figure 1 for the SHO and equations 1-6 piezoelectric hysteresis loops (*see Supplemental Materials S3*), respectively. Since all the operators in the empirical function are differentiable, the model can be trained end-to-end using backpropagation. We optimize the model to minimize the mean-squared reconstruction error (MSE) between the input spectra and the fit results.

Our model combines 1D convolutional layers with pooling for computational efficiency. These layers use learnable weight kernels with a set stride to consider the ordinal nature of the voltage sequence applied to the ferroelectric material. After dimensionality reduction, fully connected layers predict fitting parameters. The model's performance is robust to hyperparameter variations. An effective architectural choice was adding a residual layer to pass information directly from the first set of convolution layers to dense layers, aiding in capturing both general spectroscopic patterns and finer resonance details.

3.1 Accelerating the SHO fitting process

Our compact DNN for SHO fitting (Figure 1) has just 7254 parameters, in contrast to ResNet34's 21.5 million parameters[31], enabling fast training and FPGA deployment. We split the data 70/30 for training/testing to ensure our model was not overfitting. Unlike typical DNNs, our DNN uses physics constraints, leading to a non-smooth loss landscape, which is challenging for first-order optimizers like SGD. AdaHessian[32], a second-order optimizer, addresses this by approximating the second derivative, enabling adaptive learning rates. We train with AdaHessian for 5 epochs (batch size 200, initial lr 0.1) on a Tesla P100-PCIe GPU in Google Colab Pro, optimizing based on MSE between

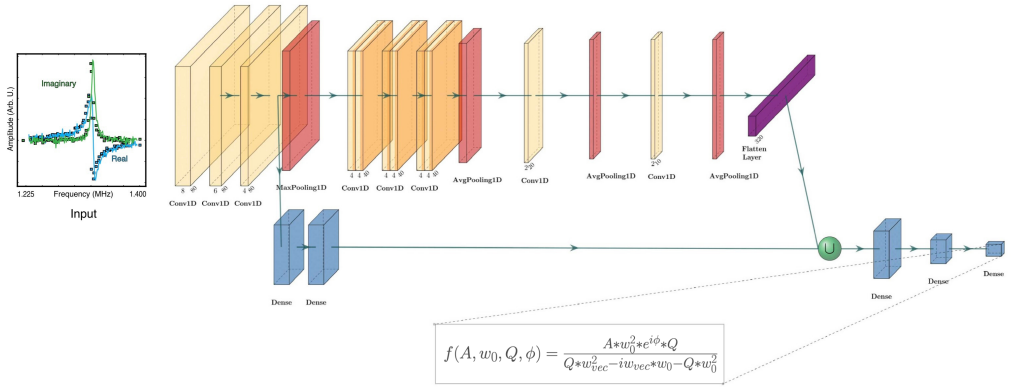

Figure 1: Deep Neural Network architectures for fitting raw data to a simple harmonic oscillator model

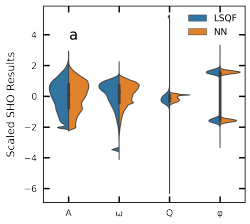

Figure 2: Distributions of predicted parameters.

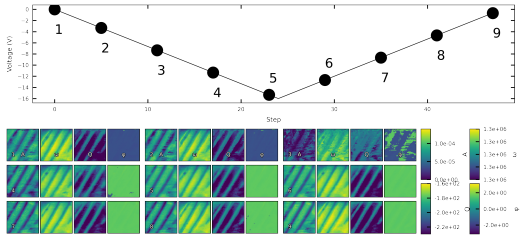

Figure 3: Band-excitation waveform with switching voltage comparing maps of four parameters (i-q) predicted by NN and LSQF.

raw data and responses. In BE-PFM, we focus on the accuracy of empirical fit parameters. LSQF and DNN show similar parameter distributions, but DNN's are consistently narrower.

Comparing amplitude, phase, resonance frequency, and quality factor maps, LSQF exhibits outliers, while the DNN provides more realistic results. In conclusion, our DNN matches or surpasses LSQF, offering practical benefits like faster fitting on GPUs (5 min vs. 18 min). (*Supplemental Materials S4*).

3.2 Performance of SHO fitting with noisy data

Our neural network's robustness was tested with SHO data corrupted by varying levels of additive random noise (noise factors: 0.0, 2.0, and 7.0) applied to both real and imaginary SHO spectra. Figure 5 compares our physics-constrained neural network (NN) with conventional least squares fitting (LSQF) on noisy synthetic data. Four key SHO fit parameters (amplitude, phase, resonance frequency, and quality factor) are analyzed across different noise factors. The neural network consistently produces more stable and reasonable parameter distributions than LSQF, which exhibits outliers, demonstrating noise-resilient fitting. At low noise (0.0), our neural network and LSQF perform well, providing accurate fit parameters. However, as noise increases (2.0 and 7.0), LSQF struggles, resulting in poor-quality fits and unrealistic parameters due to noise sensitivity and local optima convergence. Conversely, our physics-constrained neural network maintains stability even at high noise levels (7.0), reflecting its built-in regularization. While fit quality decreases with noise, the parameters remain consistent, unlike LSQF. Our network avoids strict assumptions and handles noisy signals, maintaining interpretability and physical results.

3.3 Accelerating the hysteresis loops fitting

To ensure our DNN's robustness beyond simple tasks, we tested it on complex piezoelectric hysteresis loop fitting involving a 9-parameter empirical function[33] (*see the function in Supplemental Material*

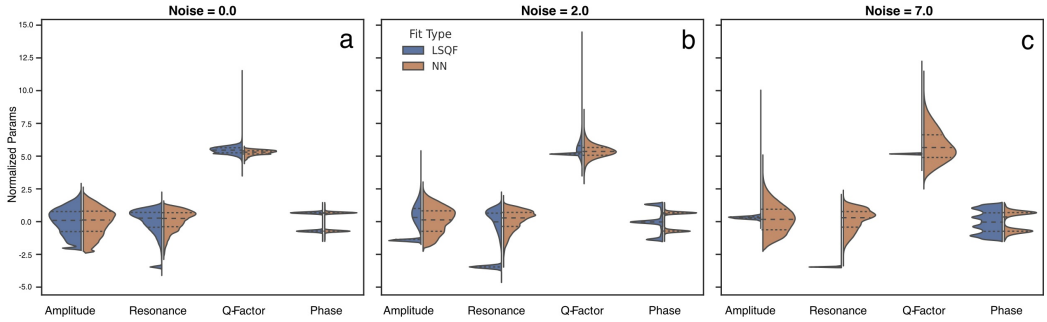

Figure 5: Violin plots comparing distributions of SHO fit parameters from the physics-constrained neural network (NN) and least squares fitting (LSQF) on noisy synthetic data. Plots are shown for amplitude (A), phase (ϕ), resonance frequency (ω), and quality factor (q) at different noise factors. The neural network maintains tighter, more reasonable distributions than LSQF, which exhibits outliers. This demonstrates the neural network's robustness to noise during fitting.

S3). This task poses challenges due to increased parameter complexity, unpredictable initial values, and parameter sensitivity. We compared our DNN with LSQF (current practice) in predicting the 9 parameters. We trained our model unsupervised to minimize MSE between observed loops and DNN-predicted fits. Optimization was crucial due to the complex function.

ADAM and AdaHessian optimization methods did not yield satisfactory results, but a Trust Region (TR) conjugate gradient optimizer outperformed them, achieving a much lower MSE in just 300 seconds. We visually compared DNN and LSQF reconstructions and used violin plots to compare NN and LSQF parameter distributions, demonstrating the DNN's superior performance (*see Figure 4 and Supplemental Materials S3 for more information*).

4 Conclusion

These tests reveal promising capabilities to robustly analyze experimental data with significant noise using our physics-constrained deep learning approach. Our method provides a pathway to enhance the value of spectroscopic techniques on systems previously considered intractable. The developed unsupervised deep neural network greatly accelerates the data analysis process in scanning probe spectroscopy, surpassing traditional methods like least squares fitting in both speed and accuracy. This advancement is crucial in scientific research where real-time feedback is essential. By efficiently handling high-dimensional data, even in noisy conditions, this approach paves the way for more creative and dynamic experimentation, particularly in fields involving complex material dynamics and nanoscale phenomena.

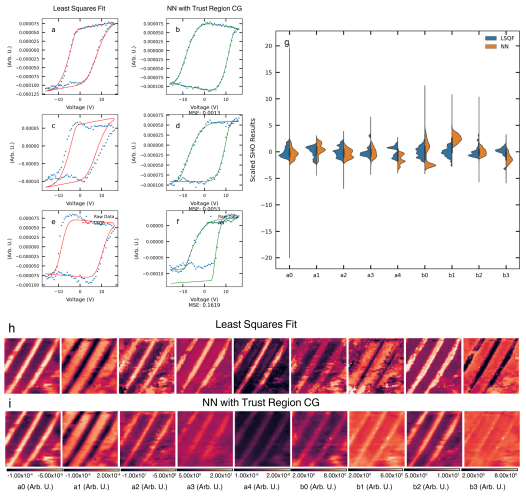

Figure 4: Piezoelectric hysteresis loops fitting results of DNN compared with LSQF method's results. a,c,e Best, median, and worst predictions of LSQF method. b,d,f Best, median, and worst predictions of a neural network trained with Trust Region CG. g Distributions of predicted parameters. h Color maps of the signal of parameters resulted from the LSQF method. i Color maps of the signal of parameters resulting from the neural network.

References

[1] F. (Feng) Tao and P. A. Crozier. Atomic-scale observations of catalyst structures under reaction conditions and during catalysis. *Chem. Rev.*, 116:3487–3539, 2016.

[2] X.-K. Wei, N. Domingo, Y. Sun, N. Balke, R. E. Dunin-Borkowski, and J. Mayer. Progress on emerging ferroelectric materials for energy harvesting, storage and conversion. *Adv. Energy Mater.*, 12:2201199, 2022.

[3] Z. Gu, S. Pandya, A. Samanta, S. Liu, G. Xiao, C. J. G. Meyers, A. R. Damodaran, H. Barak, A. Dasgupta, S. Saremi, A. Polemi, L. Wu, A. A. Podpirka, A. Will-Cole, C. J. Hawley, P. K. Davies, R. A. York, I. Grinberg, L. W. Martin, and J. E. Spanier. Resonant domain-wall-enhanced tunable microwave ferroelectrics. *Nature*, 560:622–627, 2018.

[4] S. R. Spurgeon, C. Ophus, L. Jones, A. Petford-Long, S. V. Kalinin, M. J. Olszta, R. E. Dunin-Borkowski, N. Salmon, K. Hattar, W.-C. D. Yang, R. Sharma, Y. Du, A. Chiaramonti, H. Zheng, E. C. Buck, L. Kovarik, R. L. Penn, D. Li, X. Zhang, M. Murayama, and M. L. Taheri. Towards data-driven next-generation transmission electron microscopy. *Nat. Mater.*, 20:274–279, 2021.

[5] S. V. Kalinin, B. G. Sumpter, and R. K. Archibald. Big–deep–smart data in imaging for guiding materials design. *Nat. Mater.*, 14:973–980, 2015.

[6] S. V. Kalinin, E. Strelcov, A. Belianinov, S. Somnath, R. K. Vasudevan, E. J. Lingerfelt, R. K. Archibald, C. Chen, R. Proksch, N. Laanait, and S. Jesse. Big, deep, and smart data in scanning probe microscopy. *ACS Nano*, 10:9068–9086, 2016.

[7] Tem frontiers.

[8] R. Giridharagopal, J. T. Precht, S. Jariwala, L. Collins, S. Jesse, S. V. Kalinin, and D. S. Ginger. Time-resolved electrical scanning probe microscopy of layered perovskites reveals spatial variations in photoinduced ionic and electronic carrier motion. *ACS Nano*, 13:2812–2821, 2019.

[9] J. C. Agar, B. Naul, S. Pandya, S. van der Walt, J. Maher, Y. Ren, L.-Q. Chen, S. V. Kalinin, R. K. Vasudevan, Y. Cao, J. S. Bloom, and L. W. Martin. Revealing ferroelectric switching character using deep recurrent neural networks. *Nat. Commun.*, 10:4809, 2019.

[10] V. G. Gisbert, S. Benaglia, M. R. Uhlig, R. Proksch, and R. Garcia. High-speed nanomechanical mapping of the early stages of collagen growth by bimodal force microscopy. *ACS Nano*, 15:1850–1857, 2021.

[11] R. Cohn and E. Holm. Unsupervised machine learning via transfer learning and k-means clustering to classify materials image data. *Integrating Materials and Manufacturing Innovation*, 10:231–244, 2021.

[12] S. Chen, S. Yuan, Z. Hou, Y. Tang, J. Zhang, T. Wang, K. Li, W. Zhao, X. Liu, L. Chen, L. W. Martin, and Z. Chen. Recent progress on topological structures in ferroic thin films and heterostructures. *Adv. Mater.*, 33:e2000857, 2021.

[13] Api reference. scikit-learn.

[14] Z. Wang, Z. Sun, H. Yin, X. Liu, J. Wang, H. Zhao, C. H. Pang, T. Wu, S. Li, Z. Yin, and X.-F. Yu. Data-driven materials innovation and applications. *Adv. Mater.*, page e2104113, 2022.

[15] J. C. Agar, Y. Cao, B. Naul, S. Pandya, S. van der Walt, A. I. Luo, J. T. Maher, N. Balke, S. Jesse, S. V. Kalinin, R. K. Vasudevan, and L. W. Martin. Machine detection of enhanced electromechanical energy conversion in pbzr0.2ti0.8o3 thin films. *Advanced Materials*, 30:e1800701, 2018.

[16] S. M. Ng and H. Yazid. Analysis of feature representation in dictionary learning and sparse coding algorithms for low-resolution image. *IOP Conf. Ser.: Mater. Sci. Eng.*, 864:012139, 2020.

[17] R. Kannan, A. V. Ievlev, N. Laanait, M. A. Ziatdinov, R. K. Vasudevan, S. Jesse, and S. V. Kalinin. Deep data analysis via physically constrained linear unmixing: universal framework, domain examples, and a community-wide platform. *Adv Struct Chem Imaging*, 4:6, 2018.

[18] S. Qin, Y. Guo, A. T. Kaliyev, and J. C. Agar. Why it is unfortunate that "better faster and less biased linear machine learning models work so well in electromechanical switching of ferroelectric thin films (under review). *Adv. Mater.*, 2022.

[19] S. V. Kalinin, J. J. Steffes, Y. Liu, B. D. Huey, M. Ziatdinov, K. K. Vasudevan, D. Kim, Y. Sharma, M. Ahmadi, and M. Ziatdinov. Disentangling ferroelectric domain wall geometries and pathways in dynamic piezoresponse force microscopy via unsupervised machine learning. *Nanotechnology*, 33:055707, 2021.

[20] Y. Liu, R. K. Vasudevan, K. K. Kelley, D. Kim, Y. Sharma, M. Ahmadi, S. V. Kalinin, and M. Ziatdinov. Decoding the shift-invariant data: applications for band-excitation scanning probe microscopy. *Mach. Learn.: Sci. Technol.*, 2:045028, 2021.

[21] C. M. Pate, J. L. Hart, and M. L. Taheri. Rapideels: machine learning for denoising and classification in rapid acquisition electron energy loss spectroscopy. *Sci. Rep.*, 11:19515, 2021.

[22] V. Böhm and U. Seljak. Probabilistic auto-encoder. *arXiv [cs.LG]*, 2020.

[23] G. P. P. Pun, R. Batra, R. Ramprasad, and Y. Mishin. Physically informed artificial neural networks for atomistic modeling of materials. *Nat. Commun.*, 10:2339, 2019.

[24] M. Raissi, P. Perdikaris, and G. E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.*, 378:686–707, 2019.

[25] A. T. Mohan, N. Lubbers, D. Livescu, and M. Chertkov. Embedding hard physical constraints in neural network coarse-graining of 3d turbulence. *arXiv [physics.comp-ph]*, 2020.

[26] Q. He, P. Stinis, and A. M. Tartakovsky. Physics-constrained deep neural network method for estimating parameters in a redox flow battery. *J. Power Sources*, 528:231147, 2022.

[27] S. L. Brunton, J. L. Proctor, and J. N. Kutz. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proc. Natl. Acad. Sci. U. S. A.*, 113:3932–3937, 2016.

[28] J. C. Agar, A. R. Damodaran, S. Pandya, S. van der Walt, J. T. Maher, N. Balke, S. Saremi, Q. Li, J. Kim, M. R. McCarter, L. R. Dedon, T. Angsten, N. Balke, and S. V. Kalinin L. W. Martin S. Jesse, M. Asta. Three-state ferroelastic switching and large electromechanical responses in pbtio3 thin films. *Adv. Mater.*, 29:1702069, 2017.

[29] scipy.optimize.leastsq — scipy v1.9.0 manual.

[30] N. Borodinov, S. Neumayer, S. V. Kalinin, O. S. Ovchinnikova, R. K. Vasudevan, and S. Jesse. Deep neural networks for understanding noisy data applied to physical property extraction in scanning probe microscopy. *npj Computational Materials*, 5:25, 2019.

[31] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.

[32] Z. Yao, A. Gholami, S. Shen, M. Mustafa, K. Keutzer, and M. W. Mahoney. Adahessian: An adaptive second order optimizer for machine learning, 2020.

[33] S. Jesse, H. N. Lee, and S. V. Kalinin. Quantitative mapping of switching behavior in piezoresponse force microscopy. *Rev. Sci. Instrum.*, 77:073702, 2006.